

Developing and Deploying AI/ML Applications on Red Hat OpenShift AI (AI267)

Duration: 3 Days Course Code: AI267

Overview:

An introduction to developing and deploying AI/ML applications on Red Hat OpenShift AI.

Developing and Deploying AI/ML Applications on Red Hat OpenShift AI (AI267) provides students with the fundamental knowledge about using Red Hat OpenShift for developing and deploying AI/ML applications.

Organizations collect and store vast amounts of information from multiple sources. With Red Hat OpenShift AI, organizations have a platform ready to analyze data, visualize trends and patterns, and predict future business outcomes by using machine learning and artificial intelligence algorithms.

This course helps students build core skills for using Red Hat OpenShift AI to train, develop and deploy machine learning models through hands-on experience.

You will understand the foundations of the Red Hat OpenShift AI architecture. You will be able to install Red Hat OpenShift AI, manage resource allocations, update components and manage users and their permissions.

You will also be able to train, deploy and serve models, including how to use Red Hat OpenShift AI to apply best practices in machine learning and data science.

Finally you will be able to define and set up data science pipelines with Red Hat OpenShift AI.

This course is based on Red Hat OpenShift 4.16, and Red Hat OpenShift AI 2.13.

Note: Starting January 1, 2026, Red Hat introduces RHLS-Course — a flexible subscription model now included with this catalog offering. This replaces the previous direct virtual class enrollment from Global Knowledge.

When you purchase this item, you'll receive an RHLS subscription at the course level, giving you the freedom to choose the schedule that works best and self-enroll in your selected class.

Your RHLS subscription includes:

- One live, instructor-led virtual session
- 12 months of self-paced learning access
- One certification exam with a free retake

Onsite Classroom-based sessions and closed course options remain unchanged.

Updated Jan2026

Company Events

These events can be delivered exclusively for your company at our locations or yours, specifically for your delegates and your needs. The Company Events can be tailored or standard course deliveries.

Target Audience:

- Data scientists and AI practitioners who want to use Red Hat OpenShift AI to build and train ML models
- Developers who want to build and integrate AI/ML enabled applications
- MLOps engineers responsible for installing, configuring, deploying, and monitoring AI/ML applications on Red Hat OpenShift AI

Objectives:

■ After this course participants should be able to:	■ Use Training Models
■ Understand the basis of Red Hat OpenShift AI	■ Enhance Model Training with RHOAI
■ Handle Data Science Projects	■ Get an introduction to Model Serving
■ Use Jupyter Notebooks	■ Use Model Serving in Red Hat OpenShift AI
■ Manage Red Hat OpenShift AI Installation	■ Get an introduction to Data Science Pipelines
■ Handle Users and Resources Management	■ Work with Pipelines
■ Custom Notebook Images	■ Control Pipelines and Experiments
■ Get an introduction to Machine Learning	

Prerequisites:

- Experience with Git is required
- Experience in Python development is required, or completion of the Python Programming with Red Hat (AD141) course
- Experience in Red Hat OpenShift is required, or completion of the Red Hat OpenShift Developer II: Building and Deploying Cloud-native Applications (DO288) course
- Basic experience in the AI, data science, and machine learning fields is recommended

Take Red Hat free assessment to gauge whether this offering is the best fit for your skills [Red Hat Skills Assessment](#)

Testing and Certification

Red Hat Certified Specialist in OpenShift AI Exam (EX267)

Follow-on-Courses:

None

Content:

Introduction to Red Hat OpenShift AI	Custom Notebook Images	Model Serving in Red Hat OpenShift AI
Identify the main features of Red Hat OpenShift AI, and describe the architecture and components of Red Hat AI.	Creating custom notebook images, and importing a custom notebook through the Red Hat OpenShift AI dashboard	Serve trained machine learning models with OpenShift AI
Data Science Projects	Introduction to Machine Learning	Custom Model Servers
Organize code and configuration by using data science projects, workbenches, and data connections	Describe basic machine learning concepts, different types of machine learning, and machine learning workflows	Deploy and serve machine learning models by using custom model serving runtimes
Jupyter Notebooks	Training Models	Introduction to Data Science Pipelines
Use Jupyter notebooks to execute and test code interactively	Train models by using default and custom workbenches	Create, run, manage, and troubleshoot data science pipelines
Installing Red Hat OpenShift AI	Enhancing Model Training with RHOAI	Elyra Pipelines
Installing Red Hat OpenShift AI by using the web console and the CLI, and managing Red Hat OpenShift AI components	Use RHOAI to apply best practices in machine learning and data science	Creating a Data Science Pipeline with Elyra
Managing Users and Resources	Introduction to Model Serving	KubeFlow Pipelines
Managing Red Hat OpenShift AI users, and resource allocation for Workbenches	Describe the concepts and components required to export, share and serve trained machine learning models	Creating a Data Science Pipeline with KubeFlow SDK

Additional Information:

Official course book provided to participants

Further Information:

For More information, or to book your course, please call us on Head Office Tel.: +974 40316639

training@globalknowledge.qa

www.globalknowledge.com/en-qa/

Global Knowledge, Qatar Financial Center, Burj Doha, Level 21, P.O.Box 27110, West Bay, Doha, Qatar